arXiv:1704.00710v1 [cs.CV] 3 Apr 2017

Hierarchical Surface Prediction for 3D Object Reconstruction

Christian Häne, Shubham Tulsiani, Jitendra Malik
University of California, Berkeley
{chaene, shubhtuls, malik}@eecs.berkeley.edu

Abstract

Recently, Convolutional Neural Networks have shown promising results for 3D geometry prediction. They can make predictions from very little input data such as for example a single color image, depth map or a partial 3D volume. A major limitation of such approaches is that they only predict a coarse resolution voxel grid, which does not capture the surface of the objects well. We propose a general framework, called hierarchical surface prediction (HSP), which facilitates prediction of high resolution voxel grids. The main insight is that it is sufficient to predict high resolution voxels around the predicted surfaces. The exterior and interior of the objects can be represented with coarse resolution voxels. This allows us to predict significantly higher resolution voxel grids around the surface, from which triangle meshes can be extracted. Our approach is general and not dependent on a specific input type. In our experiments we show results for geometry prediction from color images, depth images and shape completion from partial voxel grids. Our analysis shows that the network is able to predict the surface more accurately than a low resolution prediction.

1. Introduction

We live in a world composed of 3D objects bounded by 2D surfaces. Fundamentally, this means that we can either represent geometry implicitly as 3D volume or explicitly as 2D mesh surface which lives within the 3D space. When working with 3D geometry in practice for many tasks a specific representation is more suited than the other.

Recently 3D prediction approaches which directly learn a function which maps from an input image to the output geometry have emerged [6, 13, 32]. The function is represented as Convolutional Neural Network (CNN) and the geometry is represented as voxel grid, which is a natural choice for CNNs. The advantage of such an approach is that it can represent arbitrary topology and allows for large shape variations. However, such approaches have a major limitation. Due to the cubic growth of the volume with increasing resolution only a coarse voxel grid is predicted. A common resolution is $32 \times 32 \times 32$. Using higher resolutions very quickly becomes computationally infeasible. Moreover, due to the 2D nature of the surface, the ratio between surface or close to the surface and non-surface voxels becomes smaller and smaller with increasing resolution. Voxels far from the surface are generally predicted correctly quite early in the training and hence do not induce any gradients which can be used to teach the system where the surface lies. Therefore, with increasing resolution also the work done on computing gradients which are uninformative increases. The underlying reason why this problem exists is

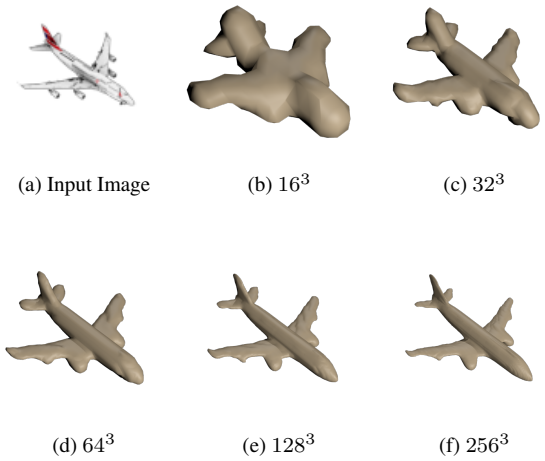

(a) Input Image (b) 16^3 (c) 32^3 (d) 64^3 (e) 128^3 (f) 256^3

Figure 1: Illustration of our method on the task of predicting geometry from a single color image. From the image our method hierarchically generates a surface with increasing resolution. In order to have sufficient information for the prediction of the higher resolution we predict at each level multiple feature channels which serve as input for the next level and at the same time allow us to generate the output of the current level. We start with a resolution of 16^3 and divide the voxel side length by two on each level reaching a final resolution of 256^3.

that current systems do not exploit the fact that surfaces are only two dimensional. Can we build a system that does?

In this paper we introduce a general framework, hierarchical surface prediction (HSP), for high resolution 3D object reconstruction which is organized around the observation that only a few of the voxels are in the vicinity of the object's surface *i.e*. the boundary between free and occupied space, while most of the voxels will be "boring", either completely inside or outside the object. The basic principle therefore is to only predict voxels around the surface, which we can now afford to do at higher resolution. The key insight to enable our method to achieve this is to change the standard prediction of free and occupied space into a three label prediction with the labels *free space*, *boundary* and *occupied space*. Furthermore, we do not directly predict high resolution voxels. Instead, we hierarchically predict small blocks of voxels from coarse to fine resolution in an octree, which we call voxel block octree. Thereby we have at each resolution the signal of the *boundary* label which tells us that the descendants of the current voxel will contain both, *free space* and *occupied space*. Therefore, only nodes containing voxels with the *boundary* label assigned need higher resolution prediction. By hierarchically predicting only those voxels we build up our octree. This effectively reduces the computational cost and hence allows us to predict significantly higher resolution voxel grids than previous approaches. In our experiments we predict voxels at a resolution of $256 \times 256 \times 256$ (*c.f*. Fig. 1).

One of the fundamental questions which arises once higher resolution predictions are feasible, is whether or not a higher resolution prediction leads to more accurate results. In order to analyze this we compare our high resolution predictions to upsampled low resolution predictions in our quantitative evaluation and determine that our method outperforms these baselines. Our results are not only quantitatively more accurate, they are also qualitatively more detailed and have a higher surface quality.

The remainder of the paper is organized as follows. In Sec. 2 we discuss the related work. We then introduce our framework in Sec. 3. Quantitative and qualitative evaluations are done in Sec. 4 and eventually we draw the conclusions in Sec. 5.

2. Related Work

Traditionally dense 3D reconstruction from images has been done using a large collection of images. Geometry is extracted by dense matching or direct minimization of reprojection errors. The common methods can be broadly grouped as implicit volumetric reconstruction [7, 18, 19, 22, 35] or explicit mesh based [9, 12] approaches. All these approaches facilitate a reconstruction of arbitrary geometry. However, in general a lot of input data is required to constrain the geometry enough. In difficult cases it is not always possible to recover the geometry of an object with multi-view stereo matching. This can be due to challenging materials such as transparent and reflective surfaces or lack of images from all around an object. For such cases object shape priors have been proposed [8, 14, 33].

All the methods mentioned above in general use multiple images as input data. Approaches which are primarily designed to reconstruct 3D geometry from a single color image were also studied. Vetter *et al*. [1] propose to build a morphable shape model for faces, which allows for the reconstruction of faces from just a single image. In order to build such a model a lot of manual interaction and high quality scanning of faces is required. Other works propose to learn a shape model from silhouettes [3] or silhouettes and keypoints [16].

Recently, CNNs have also been used to predict geometry from single images. Choy *et al*. [6] propose to use an encoder decoder architecture in a recurrent network to facilitate prediction of coarse resolution voxel grids from a single or multiple color images. An alternative approach [13] proposes to first train an autoencoder on the coarse resolution voxel grids and then train a regressor which predicts the shape code from the input image. It has also been shown that CNNs can be leveraged to predict alternative primitive based representations [31]. While these approaches rely on ground truth shapes, CNNs can also be trained to predict 3D shapes using weaker supervision [26, 32] for example image silhouettes.

Predicting geometry from color images has also been addressed in a 2.5D setting where the goal is, given a color image to predict a depth map [11, 21, 28]. These approaches can generally produce high resolution output but do only capture the geometry from a single viewpoint. Tatarchenko *et al*. [30] propose to predict a collection of RGBD (color and depth) images from a single color image and fuse them in a post-processing step.

The volumetric approach has the major limitation that using high resolution voxel grids is computationally demanding and memory intensive. Multi-view approaches alleviate this, by the use of data adaptive discretization of the volume with Delaunay tedrahedrization [20], voxel block hashing [24] or octrees [5, 29]. Our work is inspired by these earlier works on octrees. A crucial difference is that in these works the structure of the octree can be determined from the input geometry but in our case we are looking at a more general problem where the structure of the tree needs to be predicted. Similarly, a very recent work [27] which proposes to use octrees in a CNN, also assumes that the structure of the octree is given as input. We predict the structure of the tree together with its content.

Our hierarchical surface prediction framework also relates to coarse-to-fine approaches in optical flow computation. Similar to volumetric 3D reconstruction also dense

optical flow is computationally demanding due the the large label space. By computing optical flow hierarchically on a spatial pyramid, e.g. [2], real-time computation of optical flow is feasible [34]. The benefits of using a coarse-to-fine approach in learning based optical flow prediction has also recently been pointed out [25, 15].

3. Formulation

In this Section we describe our hierarchical surface prediction framework, which allows for high resolution voxel grid predictions. We first briefly discuss how the state-of-the-art, coarse resolution baseline works and then introduce the voxel block octree structure which we are predicting.

3.1. Voxel Prediction

The basic voxel prediction framework is adapted from [6, 13]. We consider an encoder/decoder architecture which takes an input $\mathcal{I}$, which in our experiments is either a color image, depth image or a partial low resolution voxel grid. A convolutional encoder encodes the input to a feature vector or shape code $\mathcal{C}$ which, in all our experiments, is a 128 dimensional vector. An up-convolutional decoder then decodes $\mathcal{C}$ into a predicted voxel grid $\mathcal{V}$. A labeling of the voxel space into free and occupied space can be determined using a suitable threshold τ (*c.f.* Sec. 4.4) or alternatively a mesh can be extracted using marching cubes [23] directly on the predicted voxel occupancies at an iso value σ.

The main problem which prevents us from directly utilizing this formulation with high resolutions for $\mathcal{V}$ is that each voxel, even if it is far from the surface, needs to be represented and a prediction is made for it. Given that the surface area grows quadratically and the volume cubically with respect to edge division, the ratio of surface to non-surface voxels becomes smaller with increasing resolution.

3.2. Voxel Block Octree

In our hierarchical surface prediction method, we propose to predict a data structure with an up-convolutional decoder architecture, which we call 'voxel block octree'. It is inspired from octree formulations [5, 29] used in traditional multi-view reconstruction approaches. The key insight which allows us to use such a data structure in a prediction framework is to extend the standard two label formulation to a three label formulation with labels *inside*, *boundary* and *outside*. As we will see later our data structure allows us to generate a complete voxel grid at high resolution while only making predictions around the surface. This leads to an approach which facilitates efficient training of an encoder/decoder architecture end-to-end.

An octree is a tree data structure which is used to partition the 3D space. The root node describes the cube of interest in the 3D space. Each internal node has up to 8 child nodes which describe the subdivision of the current node's cube into the eight octants. Note that we slightly deviate from the standard definition of the octree where either none or all the 8 child nodes are present.

We consider a tree with levels $\ell \in \{1, \dots, L\}$. Each node of the tree contains a 'voxel block' – a subdivision of the node's 3D space into a voxel grid of size b^3. Each voxel in the voxel block at tree levels $\ell < L$ contains classifier responses for the three labels *occupied space*, *boundary* and *free space*. The nodes in layers $\ell \in \{1, \dots, L-1\}$, therefore, contain voxel blocks $\mathcal{B}^{\ell,s} \in [0,1]^{b^3 \times 3}$ where the index s describes the position of the voxel block with respect to the global voxel grid. In the lowest level L we do not need the boundary label, therefore it is composed of nodes which either contain a voxel block $\mathcal{B}^{\ell,s} \in [0,1]^{b^3}$, where each element of the block describes the classifier response for the binary classification into free or occupied space. Note that each child node describes a space which has a cube side length of only half of the current node. By keeping the voxel block resolution fixed at b^3 we also divide the voxel side length by a factor of two. This means that the prediction in the child nodes are of higher resolution. The voxel block resolution b is chosen such that it is big enough for efficient prediction using an up-convolutional decoder network and at the same time small enough such that local predictions around the surface can be made. In our experiments we use voxel blocks with $b = 16$ and a tree with depth $L = 5$. A visualization of the voxel block octree is given in Fig. 3.

The predictions stored at any given level of the tree might be sparse. In order to reconstruct a complete model in high resolution, we upsample all the voxels which have not been predicted on the highest resolution from the closest predicted resolution. As we will see later, the voxels which get upsampled are in the inside and outside of the object not directly next to the boundary. Therefore the resolution of the actual surface remains high. In order to be able to extract a smooth high quality surface using marching cubes it is crucial that all the voxels which are close to the surface are predicted at high resolution and the predicted values are smooth. Therefore we aim to always evaluate the highest resolution around the boundary. At this point we would like to note that we can also extract a model at intermediate resolution by considering the boundary label as occupied space. This choice is consistent with the choice of considering any voxel as occupied space which intersects the ground truth mesh during voxelization (*c.f.* Sec. 4.1).

3.3. Network Architecture

In the previous section we introduced our voxel block octree data structure. We will now describe the encoder / decoder architecture which we are using to predict voxel block octrees using CNNs. An overview of our architecture is depicted in Fig. 2.

The encoder part is identical to the one of the basic voxel

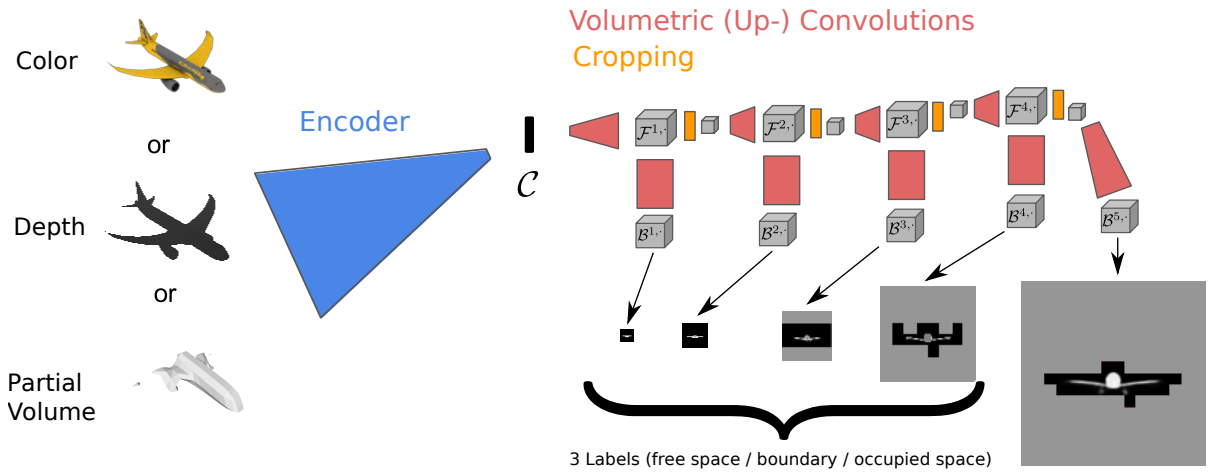

Figure 2: Overview of our system

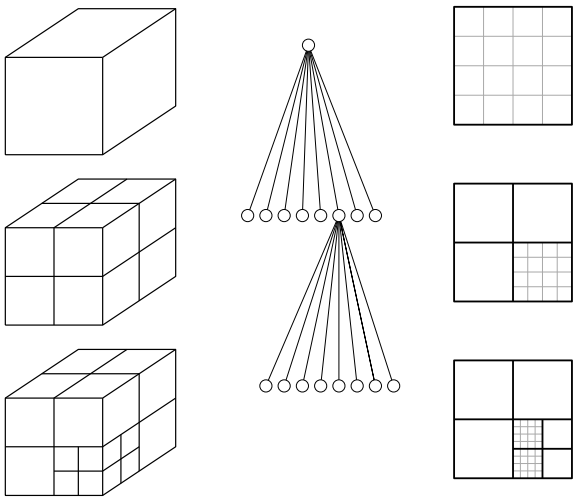

Figure 3: Visualization of the voxel block octree. The (left) part depicts the space subdivision which is induced by the octree in the (middle). The (right) side is a 2D visualization of parts of the voxel blocks with a voxel block size of $b = 4$. Note only a subset of the leaf nodes contains a voxel block, the others are empty.

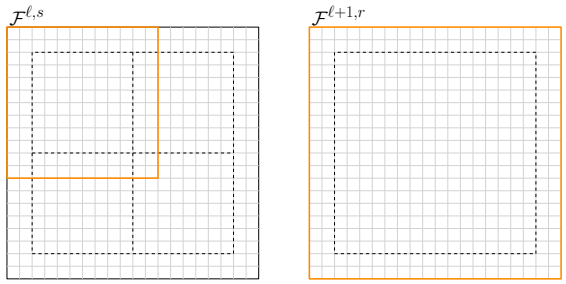

(a) Cropping (b) Upsampling

Figure 4: 2D visualization of the cropping and upsampling module. The cropping module (left) crops out the part of the feature block centered around the child node's octant. The upsampling module (right) then upsamples the feature map using up-convolutional layers to a new feature block with higher spatial resolution. The dashed lines indicate the size of the output blocks and the octants.

prediction pipeline from Sec. 3.1, *i.e.* the input $\mathcal{I}$ gets first transformed into a shape code $\mathcal{C}$ by a convolutional encoder network. An up-convolutional also called de-convolutional decoder architecture [10, 36] predicts the voxel block octree. In order to be able to predict the blocks $\mathcal{B}^{\ell,s}$ and at the same time also being able to capture the information which allows for predictions of the higher resolution levels $\ell + 1, \ldots$, we introduce feature blocks $\mathcal{F}^{\ell,s} \in \mathbb{R}^{(b+2p)^3 \times c}$. The spatial extent of the feature blocks is bigger or equal to the ones of the voxels blocks to allow for a padding $p \geq 0$. As we will see later the padding allows for an overlap when predicting neighbouring octants which leads to smoother results. The fourth dimension is the number of feature channels which can be chosen depending on the application. In our experiments we use $c = 32$ and $p = 2$.

The most important part for the understanding of our decoder architecture is how we predict level $\ell + 1$ given level ℓ. In the following we will discuss this procedure in detail. It consist of three basic steps.

1. Feature Cropping
2. Upsampling

(a) Input (b) No Overlap (c) 2 Voxels Overlap

Figure 5: Influence of overlap on the task of geometry prediction from a single depth map. Without overlap the predictions are much less smooth and the grid pattern induced by the individual voxel blocks is much more visible.

3. Output Generation

Feature Cropping. At this point we assume that we have given the feature block $\mathcal{F}^{\ell,s}$. The goal is to generate the output for the child node corresponding to a specific octant $\mathcal{O}$. The feature block $\mathcal{F}^{\ell,s}$ contains information to generate the output for all the child nodes. In order to only process the information relevant for the prediction of the child node corresponding to $\mathcal{O}$ we extract a $((b/2+2p)^3 \times c)$ region out of the four dimensional tensor $\mathcal{F}^{\ell,s}$ spatially centered around $\mathcal{O}$. An illustration of this process is given in Fig. 4. If neighboring octants are processed, the extracted feature channels will have some overlap, which helps to generate a smoother output (see Fig. 5).

Upsampling. The upsampling module takes input from the cropping module and predicts a new b^3 feature block $\mathcal{F}^{\ell+1,r}$ via up-convolutional and convolutional layers.

Output Generation. The output network takes the prediction of the feature block $\mathcal{F}^{\ell+1,r}$ from the upsampling module and generates the voxel block $\mathcal{B}^{\ell+1,r}$. This is done using a sequence of convolutional layers. The supervision is given in form of the three ground truth labels for the voxel block $\mathcal{B}^{\ell+1,r}$, *i.e.* there is supervision for each level of the tree. Once the output is generated the child nodes for the next level get generated. The decision on whether to add a child node and hence a higher resolution prediction is based on the boundary prediction in the corresponding octant of the voxel block $\mathcal{B}^{\ell+1,r}$. We compute the maximum boundary prediction response of the corresponding octant $\mathcal{O}'$

$$C^{\ell+1,r}_{\mathcal{O}'} = \max_{i,j,k \in \mathcal{O}'} \mathcal{B}^{\ell+1,r}_{i,j,k,2} \ , \tag{1}$$

with label 2 the *boundary* label. Output for the child node is generated if $C^{\ell+1,r}_{\mathcal{O}'}$ is above a small threshold γ. The intuition behind this choice is that as soon as there is some evidence that the surface is going through a specific block we should predict a higher resolution output. On the other hand if the prediction on a specific level is very certain that there is no boundary within that subtree, a higher resolution prediction is not necessary. For the classes aeroplane, chair and car in practice only around 5 to 10% of the voxels are predicted at high resolution, in average (*c.f.* Fig. 6).

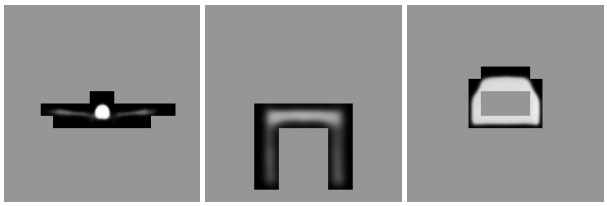

Figure 6: Responses at the highest resolution, the gray areas mean not predicted at that resolution, (left) slice through airplane, (middle) slice through front legs of a chair, (right) slice through a car.

In the first level of the tree, which predicts the root node from the shape code $\mathcal{C}$ a small decoding network directly predicts the first feature block without a cropping module in between. Likewise, in the deepest level of the tree no explicit feature block is needed. Therefore the output is directly generated from the cropped features of the previous level. Also note that the output and upsampling modules have their individual filters at each level of the tree.

3.4. Efficient Training with Subsampling

As we have seen above, only predicting voxels on the boundary at each level largely reduces the complexity of the prediction task. However, during training, a complete prediction of the tree is still too expensive. Especially in the beginning of training when the boundary predictions are very blurry and inaccurate, this would lead to many tree evaluations. Moreover, most of these evaluations would be at a location far away from the boundary, not actually giving any training signal for the learning of that specific level.

Therefore, we propose to utilize a subsampling of the child nodes during training. The tree gets traversed in a depth first manner. Each time a child node needs to be visited this only happens with a certain probability p. We start training with a small probability p in order to mostly train the encoder and the first level of the voxel block octree in the beginning. As soon as the decrease of the validation loss starts to be slow we increase p to ensure that all the levels of the tree get eventually trained. The procedure of increasing p is repeated until the validation loss does not further decrease. Note that this is just one specific choice for training the network. The main focus of our experiments was not on finding the most efficient training procedure.

In order to be able to train with mini-batches we sum up all the gradients during traversal of the tree and only do a gradient step as soon as we have done a forward and backward traversal of the subsampled tree for the training ex-

amples of the whole mini-batch. As loss function we use Cross-Entropy.

4. Experiments

For our evaluation we use the synthetic dataset ShapeNetCore [4] . It is composed of Computer Aided Design (CAD) models of objects which are organized in categories. We use three categories for our evaluation, aeroplanes, chairs and cars. Our system is implemented in Torch[1] and we use Adam [17] for stochastic optimization, with a mini-batch size of 5.

4.1. Voxelization of the Ground Truth

In a prepossessing step we voxelize the ground truth data. A common approach is to consider all the voxels which intersect the ground truth mesh as occupied space and then fill in the interior by labeling all voxels which are not reachable through free sapce from the boundary as occupied space. While this works well for the commonly used 32^3 resolution it does not directly generalize to high resolutions. The reason is that the CAD models are generally not watertight meshes and with increasing resolution the risk that a hole in the mesh prevents filling the interior increases. In our case we are voxelizing the models at a resolution of 256^3. In order to have filled interiors of the objects we utilize a multi-scale approach. We first fill the interior at 32^3 resolution. Then erode the occupied space by one voxel. We then fix every high resolution voxel which falls within this eroded low resolution occupied space as occupied space and also fix every high resolution voxel which falls within the original low resolution free space as free space. Furthermore, we fix all the high resolution voxels with intersect the original mesh surface as occcupied space. To determine the ground truth label for the remaining high resolution voxels we run a graph-cut based regularization with a small smoothness term and a preference to keep the unlabeled voxels as free space.

Given the high resolution voxels at 256^3 resolution we can now build the ground truth for all the remaining levels of the pyramid. At each level of the pyramid we label the voxels which contain the boundary at highest resolution as boundary and the other labels either free or occupied space, our coarsest resolution is 16^3.

4.2. Baselines

We consider two baselines. Both of the baselines use the state-of-the-art coarse resolution voxel prediction framework from Sec. 3.1. For both baselines we use a prediction resolution of 32^3, they differ in the way the ground truth is computed. The first baseline follows the standard approach of labeling all voxels which intersect the ground truth mesh surface as occupied space and then fill in the interior. This can be achieved by downsampling our high resolution ground truth and label all low resolution voxels which contain at least one high resolution occupied space voxel as occupied space and all the other ones as free space. We call this baseline "Low Resolution Hard" (LR Hard). The other baseline uses a soft assignment for the low resolution, the label is given by the ratio of high resolution free to occupied space voxels that the low resolution voxel contains. Therefore these labels can have fractional assignments. We call this baseline "Low Resolution Soft" (LR Soft). Note that the baseline LR Soft makes use of the high resolution voxelization but the baseline LR Hard is equivalent to voxelizing at low resolution. As our goal is prediction of high resolution geometry we bilinearly upsample the raw classification output of the baselines from 32^3 to 256^3 and conduct the evaluation at high resolution.

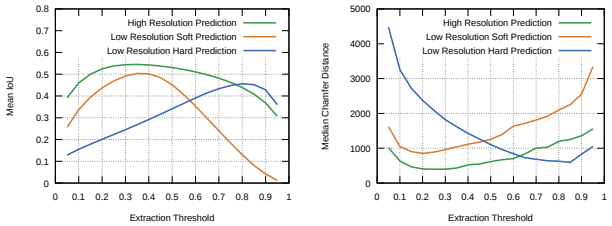

Figure 7: Evaluation on the validation set at different extraction thresholds for the class aeroplane. For IoU higher is better and for CD lower is better.

Category	HSP	LR Soft	LR Hard
Aeroplane	**0.544**	0.504	0.457
Chair	**0.361**	0.339	0.351
Car	**0.698**	0.675	0.640

Table 1: Average Intersection over Union (IoU) measure

Category	HSP	LR Soft	LR Hard
Aeroplane	**503**	870	687
Chair	3603	4181	**3546**
Car	**1513**	1677	2040

Table 2: Median Chamfer Distance measure

4.3. Input Data

To obtain the RGB/Depth images which are used as input for training and testing our CNNs, we render the CAD models in the ShapeNet dataset using Blender[2]. For each CAD model, we render 10 images from random viewpoints obtained via uniformly sampling azimuth from $[0, 360)$ degrees and elevation from $[-20, 30]$ degrees. We also use

[1] http://torch.ch/

[2] http://www.blender.org

Figure 8: Results for the category chairs. Each block from left to right, input image, our proposed HSP, LR Soft, LR Hard.

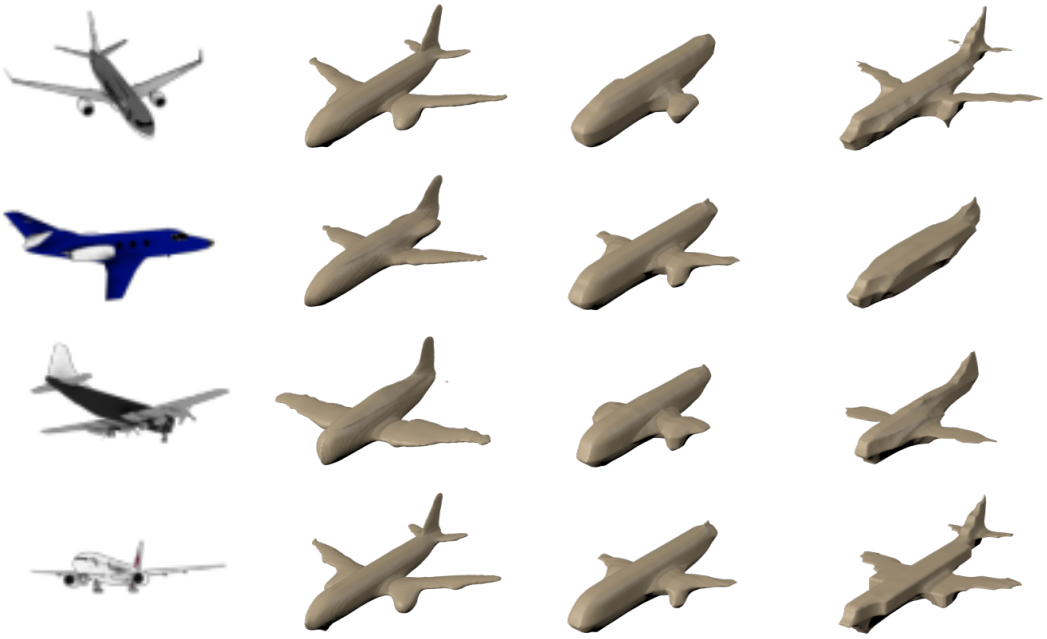

Figure 9: Results for the category aeroplanes. Each block from left to right, input image, our proposed HSP, LR Soft, LR Hard.

random lighting variations to render the RGB images.

We also use partial voxel grids as input. In practice such an input can arise when multiple depth maps of an object are available but they are all from the same side of the object. To imitate this task we use the ground truth of the LR Soft baseline and randomly zero the data in half of the voxel grid. The network then learns to predict the complete high resolution geometry from the partial low resolution input.

4.4. Quantitative Evaluation

We conducted a quantitative evaluation on the task of predicting high resolution geometry from a single RGB image. We assign each of the 3D models from the dataset randomly to the train, validation or test split with probabil-

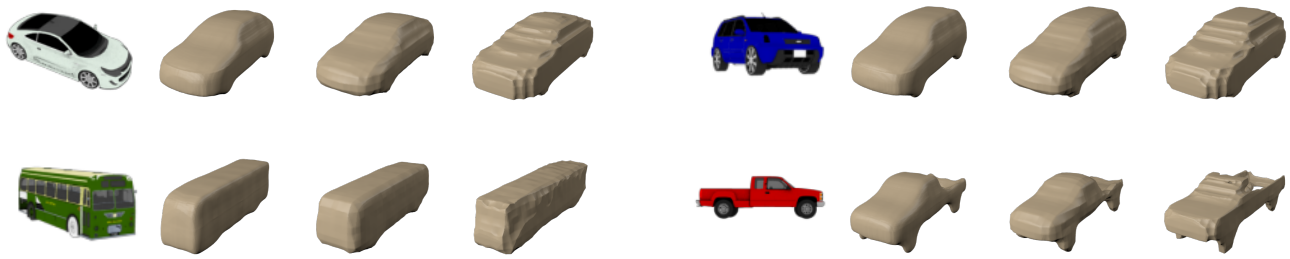

Figure 10: Results for the category cars. Each block from left to right, input image, our proposed HSP, LR Soft, LR Hard.

Figure 11: Additional applications of our method, (left) reconstruction from real world rgb images, (right) geometry prediction from partial low resolution input volumes.

ities 70%, 10% and 20%, respectively. We trained category specific networks for our method Hierarchical Surface Prediction (HSP) and both baselines for the three categories, aeroplanes, chairs and cars. As error measures we use the Intersection over Union (IoU) and the symmetric Chamfer Distance (CD). To compute the error measures we binarize the non-binary predictions with a suitable threshold. In order to determine the best threshold for each error measure and category we test an exhaustive set of thresholds using the validation set. The plots given in Fig. 7 show that the choice of a good threshold has a very drastic influence to the error measure. We compute the IoU and Chamfer distance for each of the models in the test set using the per category and per error measure best threshold determined using the validation set. The results are given in Tables 1 and 2. In all cases except for the Chamfer distance on chairs we outperform the baselines.

4.5. Qualitative Evaluation

In Figs. 8,9 and 10 we show example results on the task of geometry prediction from a single color image for the three categories. The results are predicted with the same networks which we used in the quantitative analysis above. The meshes are extracted with marching cubes [23] using the same thresholds as for the quantitative analysis. For chair the LR Hard baseline produces visually pleasant results. However, this same baseline does not achieve detailed results on the categories aeroplane and car. Due to the fractional labels the LR Soft baseline often misses thin structures. The results generated with our proposed hierarchical surface prediction have high quality surfaces for all the three categories.

We also downloaded two images of cars with a white background from the Internet using Google image search and processed them with our car prediction network, which was trained only on synthetic data. As can be seen in Fig. 11, the network is able to generalize to this type of input. Additionally, we trained a network on the category car which takes a partial low resolution voxel grid as input and predicts a high resolution model of the whole object. In practice such types of input occurs when for example multiple views of only one side of an object are given. Two results of this type of prediction are depicted in Fig. 11.

5. Conclusion and Future Work

In this work, we presented a hierarcical surface prediction framework which facilitates high resolution 3D reconstruction of objects from very little input data. We presented results on the task of predicting geometry from a single color image, a single depth image and from a partial low resolution voxel grid. Our quantitative evaluation shows that the high resolution prediction leads to more accurate results than low resolution baselines. Qualitatively, our predicted surfaces are superior to the low resolution baselines.

In this work we have demonstrated that high resolution geometry prediction is feasible therefore future work will investigate the potential to use prediction based reconstruction methods in multi-view stereo. Furthermore, for almost symmetric objects methods which explicitly take the symmetry into account are an interesting further direction which might improve the quality of prediction based approaches.

Acknowledgments

Christian Häne is receiving an "Early Postdoc.Mobility" fellowship No. 165245 from the Swiss National Science Foundation. The project received funding from the Intel/NSF VEC award IIS-153909 and NSF award IIS-1212798.

References

[1] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Conference on Computer graphics and interactive techniques (SIGGRAPH)*, 1999. 2

[2] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision (ECCV)*, 2004. 3

[3] T. J. Cashman and A. W. Fitzgibbon. What shape are dolphins? building 3d morphable models from 2d images. *Transactions on pattern analysis and machine intelligence (TPAMI)*, 2013. 2

[4] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 6

[5] J. Chen, D. Bautembach, and S. Izadi. Scalable real-time volumetric surface reconstruction. *Transactions on Graphics (TOG)*, 2013. 2, 3

[6] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 2, 3

[7] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Conference on Computer graphics and interactive techniques (SIGGRAPH)*, 1996. 2

[8] A. Dame, V. A. Prisacariu, C. Y. Ren, and I. Reid. Dense reconstruction using 3d object shape priors. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 2

[9] A. Delaunoy, E. Prados, P. G. I. Piracés, J.-P. Pons, and P. Sturm. Minimizing the multi-view stereo reprojection error for triangular surface meshes. In *British Machine Vision Conference (BMVC)*, 2008. 2

[10] A. Dosovitskiy, J. T. Springenberg, M. Tatarchenko, and T. Brox. Learning to generate chairs, tables and cars with convolutional networks. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, 2017. 4

[11] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems (NIPS)*, 2014. 2

[12] P. Gargallo, E. Prados, and P. Sturm. Minimizing the reprojection error in surface reconstruction from images. In *International Conference on Computer Vision (ICCV)*, 2007. 2

[13] R. Girdhar, D. F. Fouhey, M. Rodriguez, and A. Gupta. Learning a predictable and generative vector representation for objects. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 2, 3

[14] C. Häne, N. Savinov, and M. Pollefeys. Class specific 3d object shape priors using surface normals. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2

[15] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. *arXiv preprint arXiv:1612.01925*, 2016. 3

[16] A. Kar, S. Tulsiani, J. Carreira, and J. Malik. Category-specific object reconstruction from a single image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2

[17] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6

[18] K. Kolev, M. Klodt, T. Brox, S. Esedoglu, and D. Cremers. Continuous global optimization in multiview 3d reconstruction. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, 2007. 2

[19] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. *International Journal of Computer Vision (IJCV)*, 2000. 2

[20] P. Labatut, J.-P. Pons, and R. Keriven. Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In *International Conference on Computer Vision (ICCV)*, 2007. 2

[21] L. Ladicky, J. Shi, and M. Pollefeys. Pulling things out of perspective. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2

[22] V. Lempitsky and Y. Boykov. Global optimization for shape fitting. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007. 2

[23] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Siggraph computer graphics*, 1987. 3, 8

[24] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3d reconstruction at scale using voxel hashing. *Transactions on Graphics (TOG)*, 2013. 2

[25] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. *arXiv preprint arXiv:1611.00850*, 2016. 3

[26] D. J. Rezende, S. A. Eslami, S. Mohamed, P. Battaglia, M. Jaderberg, and N. Heess. Unsupervised learning of 3d structure from images. In *Advances In Neural Information Processing Systems (NIPS)*, 2016. 2

[27] G. Riegler, A. O. Ulusoys, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. *arXiv preprint arXiv:1611.05009*, 2016. 2

[28] A. Saxena, S. Chung, and A. Ng. Learning depth from single monocular images. In *Neural Information Processing Systems (NIPS)*, 2005. 2

[29] F. Steinbrücker, J. Sturm, and D. Cremers. Volumetric 3d mapping in real-time on a cpu. In *International Conference on Robotics and Automation (ICRA)*, 2014. 2, 3

[30] M. Tatarchenko, A. Dosovitskiy, and T. Brox. Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision (ECCV)*, 2016. 2

[31] S. Tulsiani, H. Su, L. J. Guibas, A. A. Efros, and J. Malik. Learning shape abstractions by assembling volumetric primitives. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2

[32] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *Advances in Neural Information Processing Systems*, 2016. 1, 2

[33] S. Yingze Bao, M. Chandraker, Y. Lin, and S. Savarese. Dense object reconstruction with semantic priors. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 2

[34] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Joint Pattern Recognition Symposium (DAGM)*, 2007. 3

[35] C. Zach, T. Pock, and H. Bischof. A globally optimal algorithm for robust tv-l 1 range image integration. In *International Conference on Computer Vision (ICCV)*, 2007. 2

[36] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. Deconvolutional networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 4

A. Detailed CNN Architectures

In this section we give the detailed layer configurations of our networks, including the baselines. We use the following abbreviations:

Conv	Convolution
UpConv	Up-convolution
FC	Fully connected
MP	Max Pooling
LR	Leaky ReLU with negative slope of 0.2
R	ReLU
DO	Dropout with probability 0.5
RS	Reshape
Crop	Cropping center region.
kW, kH, kD	Kernel sizes in the three dimensions
sW, sH, sD	Strides in the three dimensions
oC	Number of output feature channels
oW, oH, oD	Output sizes in the three dimensions

For both our proposed hierarchical surface prediction and the baselines we use the same encoder. The layers of the encoder for color or depth images are given in Tab. 3. The encoder for the partial low resolution grids has the same general structure. The decoder for both the baselines is given in Tab. 4. Note that both baselines use the exact same network. In one case the ground truth are hard binary labels and in the other case the ground truth are soft labels which represent the ratio of free/occupied space of the high resolution voxels, *c.f.* main text.

Our proposed hierarchical surface prediction (HSP) uses a set of modules in the decoder. The module which decodes the shape code from the bottle neck to the first feature block is given in Tab. 5. As described in the main text each transition from one level to the next is done using an upsampling module Tab. 6 and an output module Tab. 7. On the last level the output is generated by the module given in Tab. 8.

lType	kW	kH	sW	sH	oC	oW	oH
Input	-	-	-	-	1/3	128	128
Conv	3	3	1	1	8	128	128
MP + LR	2	2	2	2	8	64	64
Conv	3	3	1	1	16	64	64
MP + LR	2	2	2	2	16	32	32
Conv	3	3	1	1	32	32	32
MP + LR	2	2	2	2	32	16	16
Conv	3	3	1	1	64	16	16
MP + LR	2	2	2	2	64	8	8
Conv	3	3	1	1	128	8	8
MP + LR	2	2	2	2	128	4	4
Conv	3	3	1	1	128	4	4
MP + LR	2	2	2	2	128	2	2
Conv	3	3	1	1	128	2	2
MP + LR	2	2	2	2	128	1	1
DO	1	1	1	1	128	1	1
FC + LR	1	1	1	1	128	1	1
FC + LR	1	1	1	1	128	1	1

Table 3: Color/Depth Encoder

lType	kW	kH	kD	sW	sH	sD	oC	oW	oH	oD
FC + LR	1	1	1	1	1	1	1024	1	1	1
RS	-	-	-	-	-	-	128	2	2	2
Conv + R	3	3	3	1	1	1	128	2	2	2
UpConv + R	4	4	4	2	2	2	128	4	4	4
Conv + R	3	3	3	1	1	1	128	4	4	4
UpConv + R	4	4	4	2	2	2	128	8	8	8
Conv + R	3	3	3	1	1	1	64	8	8	8
UpConv + R	4	4	4	2	2	2	32	16	16	16
Conv + R	3	3	3	1	1	1	16	16	16	16
UpConv + R	4	4	4	2	2	2	8	32	32	32
Conv	3	3	3	1	1	1	1	32	32	32

Table 4: Baseline Decoder

lType	kW	kH	kD	sW	sH	sD	oC	oW	oH	oD
FC + LR	1	1	1	1	1	1	1024	1	1	1
RS	-	-	-	-	-	-	128	2	2	2
Conv + R	3	3	3	1	1	1	128	2	2	2
UpConv + R	4	4	4	2	2	2	128	4	4	4
Conv + R	3	3	3	1	1	1	128	4	4	4
UpConv + R	4	4	4	2	2	2	128	8	8	8
Conv + R	3	3	3	1	1	1	64	6	6	6
UpConv + R	4	4	4	2	2	2	64	12	12	12
Conv + R	3	3	3	1	1	1	32	12	12	12
UpConv + R	4	4	4	2	2	2	32	20	20	20
Conv + R	3	3	3	1	1	1	32	20	20	20

Table 5: Decoder module, bottleneck to feature block $\mathcal{F}^{1,1}$

lType	kW	kH	kD	sW	sH	sD	oC	oW	oH	oD
UpConv + R	4	4	4	2	2	2	32	20	20	20
Conv + R	3	3	3	1	1	1	32	20	20	20

Table 6: Upsampling module

lType	kW	kH	kD	sW	sH	sD	oC	oW	oH	oD
Conv + R	3	3	3	1	1	1	16	20	20	20
Conv + R	3	3	3	1	1	1	8	20	20	20
Conv + R	3	3	3	1	1	1	3	20	20	20
Crop	-	-	-	-	-	-	3	16	16	16

Table 7: Output module

lType	kW	kH	kD	sW	sH	sD	oC	oW	oH	oD
UpConv + R	4	4	4	2	2	2	16	20	20	20
Conv + R	3	3	3	1	1	1	8	20	20	20
Conv + R	3	3	3	1	1	1	1	20	20	20
Crop	-	-	-	-	-	-	1	16	16	16

Table 8: Level L output module